arXiv:1410.0001v1 [cs.IR] 30 Sep 2014

On Evaluation Validity in Music Autotagging

Fabien Gouyon, Bob L. Sturm, João Lobato Oliveira,
Nuno Hespanhol, and Thibault Langlois

Abstract

Music autotagging, an established problem in Music Information Retrieval, aims to alleviate the human cost required to manually annotate collections of recorded music with textual labels by automating the process. Many autotagging systems have been proposed and evaluated by procedures and datasets that are now standard (used in MIREX, for instance). Very little work, however, has been dedicated to determine what these evaluations really mean about an autotagging system, or the comparison of two systems, for the problem of annotating music in the real world. In this article, we are concerned with explaining the figure of merit of an autotagging system evaluated with a standard approach. Specifically, does the figure of merit, or a comparison of figures of merit, warrant a conclusion about how well autotagging systems have learned to describe music with a specific vocabulary? The main contributions of this paper are a formalization of the notion of validity in autotagging evaluation, and a method to test it in general. We demonstrate the practical use of our method in experiments with three specific state-of-the-art autotagging systems –all of which are reproducible using the linked code and data. Our experiments show for these specific systems in a simple and objective two-class task that the standard evaluation approach does not provide valid indicators of their performance.

1 Introduction

Music autotagging is an established problem in Music Information Retrieval (MIR), as witnessed by the publication of book chapters (e.g., [Bertin-Mahieux et al., 2010]), several journal articles (e.g., [Turnbull et al., 2008, Bertin-Mahieux et al., 2008, Fu et al., 2011, Miotto and Lanckriet, 2012]) and conference papers (e.g., [Miotto et al., 2010, Seyerlehner et al., 2010, Xie et al., 2011, Marques et al., 2011, Coviello et al., 2012, Nam et al., 2012]), PhD theses (e.g., [Sordo, 2012]), tutorials (ISMIR 2013). Music autotagging systems aim to annotate music audio signals with textual labels, or tags. Ultimately, such systems could alleviate the human cost required to manually annotate collections of recorded music by automating the process.

Many music autotagging systems have been proposed and evaluated by procedures and datasets that are now standard, as exemplified e.g. by six years of completed MIREX "Audio Tag Classification" task (ATC). The topic of system evaluation itself plays a increasingly critical role in the MIR community, as mentioned in the challenges highlighted in a recent Roadmap for MIR [Serra et al., 2013].

Clearly, the desire of this field of research is for an autotagging system, or any MIR system, to perform well in the real world. One step towards considering how well MIR systems work in the real world is testing their robustness to a variety of environmental conditions, such as noise, audio quality, etc. For instance, work has been dedicated to the effect of audio perturbations (e.g. adding white noise, filtering, different encodings, etc.) on the computation of low-level features such as MFCCs or chromas [Sigurdsson et al., 2006, Jensen et al., 2009, Urbano et al., 2014], and on the robustness to audio perturbations of state-of-the-art systems for beat tracking, chord recognition, and audio-to-score alignment [Gouyon et al., 2006, Mauch and Ewert, 2013].

Whereas robustness tests seek to determine how sensitive a system is to characteristics of its environment, we contend the question that needs to be addressed first is whether a system's evaluation provides us with valid conclusions about its true performance. Indeed, virtually no autotagging evaluation has addressed the question of validity [Urbano et al., 2013,Sturm, 2014b].

The main contributions of this paper are precisely a formalization of the notion of validity in autotagging evaluation, and a method to test it in general. This method is based on the consideration that if an autotagging system is pairing audio signals with tags in a meaningful way, its behavior should not be significantly affected by irrelevant perturbations of its input signals. We perform several experiments demonstrating our method for three state-of-the-art autotagging systems. We confirm in these experiments that the irrelevant perturbations we perform are "fair", i.e. they do not imply a significant covariate shift between the feature distributions of training and test data [Sugiyama et al., 2007, Quionero-Candela et al., 2009].

This article is organized as follows: In the next section, we clarify the objectives of evaluation in music autotagging research, review the standard approach to evaluation, and formalize the notion of validity in the context of evaluation of autotagging systems. Then, in Section 3, we present a method for testing the validity of autotagging evaluation, based on specifically designed perturbations of test instances, which we define as "irrelevant transformations." Section 4 describes our experiments with this method in testing the validity of the evaluation of three specific state-of-the-art autotagging systems. We summarize the article and discuss its findings in Section 5. All experiments and results in this article can be reproduced via data available on `http://www.fabiengouyon.org/`, under the "Research"

– "Data for reproducible research" menu item.

2 Music Autotagging and its Evaluation

2.1 What is autotagging?

Following [Turnbull et al., 2008], we consider music autotagging as a multi-label supervised learning problem with music audio signals as input, and where the objective is to meaningfully relate tag concepts and acoustic phenomena. Adopting the terminology of [Seyerlehner et al., 2010], we equate music autotagging to "transform[ing] an audio feature space into a semantic space, where music is described by words", and we define a music autotagging system as one that annotates, i.e., assigns tags to, recorded music. For example, if singing voice is heard in the music, a good music autotagging system should annotate it with the tag "vocals".

2.2 Current practices of music autotagging evaluation

An in-depth formalisation of evaluation in comparative experiments can be found in [Bailey, 2008], and a preliminary application of it to the specific case of evaluation in MIR in [Sturm, 2014a]. A standard approach to music autotagging evaluation is having a system annotate a set of signals, and then comparing the resulting tags to the "ground truth." Between 2008-2012, the MIREX[1] "Audio Tag Classification" task (ATC) has employed this approach to systematically and rigorously evaluate about 60 music autotagging solutions with standardized datasets. This evaluation procedure also appears in many other works, e.g. [Turnbull et al., 2008,Bertin-Mahieux et al., 2008, Miotto et al., 2010, Xie et al., 2011, Coviello et al., 2012, Nam et al., 2012].

A fundamental aspect of these evaluations is *data*. The music autotagging literature has established a variety of benchmark datasets. Several works use the datasets CAL500 [Turnbull et al., 2008], MagnaTagatune [Law et al., 2009], and the Million Song Dataset [Bertin-Mahieux et al., 2011]. Among the datasets ATC uses are MajorMiner [Mandel and Ellis, 2008] and USPOP [Berenzweig et al., 2004]. Evaluation in music autotagging typically proceeds via cross-validation experiments, as follows. A dataset of sampled audio signals is partitioned into K non-overlapping folds. This dataset is such that each signal is paired with "ground truth" tags from a given tag vocabulary. Then, K music autotagging systems are built by training on the complement of a testing dataset fold. The presence or absence of each tag from the "ground truth" is measured in the output of the system. More specifically, the following *measurements* are made: the number of true pos-

[1] `http://www.music-ir.org/mirex/wiki/MIREX_HOME`

itives, false positives, true negatives, and false negatives of each tag are counted.

Music autotagging evaluation involves computing several *figures of merit* (FoM) from these measurements. In ATC, these include quantities named "Average Tag Recall," "Average Tag Precision," "Average Tag F-Measure," the precise meanings of which are specified in the source code of MIREX.[2] The ATC figure of merit "Average Tag Recall" is defined as the mean of the K micro-averaged recalls (also called "global" recalls); the "Average Tag Precision" is defined as the mean of the K micro-averaged precisions; and the "Average Tag F-Measure" is defined as the mean harmonic mean of the K "Average Tag Precisions" and "Average Tag Recalls." Other figures of merit appear in the literature. For instance, the macro-averaged recall of a system is defined as the mean of the recalls of each tag. This is also called per-tag recall [Turnbull et al., 2008,Bertin-Mahieux et al., 2008,Miotto et al., 2010, Marques et al., 2011, Xie et al., 2011, Coviello et al., 2012, Nam et al., 2012]. Similarly, there is the macro-averaged precision, and macro-averaged F-measure.

2.3 What can one expect from evaluating an autotagging system?

Denote an autotagging system by S, which maps an input audio signal $\mathbf{x}$ to a subset $\mathcal{X}$ of a set of tags, denoted $\mathcal{T}$. A dataset is defined as an indexed set of tuples $(\mathbf{x}, \mathcal{X})$. We notate the training dataset Ψ and the testing dataset Φ.

A relatively common assumption to the design and evaluation of supervised learning systems, such as autotagging systems, is that the feature distributions of their training and test data are identical (i.i.d.) [Quionero-Candela et al., 2009]. That is, that the features in Ψ and Φ are sampled from the same distribution $\mathcal{D}$. For instance, [Marques et al., 2011] illustrate the fact that state-of-the-art autotagging systems trained on a given dataset typically fail to generalize to datasets of different origins, where the i.i.d. assumption is not respected. On the other hand, when the feature vectors of Ψ and Φ are i.i.d., one should expect the performance of S trained on Ψ to be relatively stable with respect to different sets Φ. This is for instance the case when Ψ and Φ are different folds (or combinations thereof) of the same dataset in a cross-validation procedure (see Section 2.2). One should therefore expect that S be put to use in "similar conditions" than those used for training.[3]

[2] See method `evaluateResultFold` in `https://code.google.com/p/nemadiy/source/browse/analytics/trunk/src/main/java/org/imirsel/nema/analytics/evaluation/tagsClassification/TagClassificationEvaluator.java`

[3] Note however that research in Domain Adaptation and Transfer Learning precisely address the design of systems coping with conditions different than those under which they were developed [Quionero-Candela et al., 2009,Pan and Yang, 2010,Ben-David et al.,

2.4 Validity in music autotagging evaluation

An evaluation of music autotagging systems produces measurements, from which FoM are computed and conclusions then drawn. For instance, when an FoM is significantly better for one system compared to another, then one desires that the former system is better at autotagging than the latter. Hence, a critical question to answer is whether the approach used for evaluation is *valid* for such conclusions, i.e. whether "we are really measuring what we want to measure" [Urbano et al., 2013].

More formally, denote by $\Gamma_S(t)$ the *true performance* of a system S on a tag $t \in \mathcal{T}$. (Note that $\Gamma_S(t)$ is a simplified notation for $\Gamma_{S;\Psi}(t)$, as the system is a product of the training dataset Ψ.) The true performance describes how well S is expected to perform in using t (or not) to annotate any test music audio signals (assuming i.i.d. between train and test data). Define $\Gamma_S(t) = \mathbb{E}\big[f_S(\mathbf{x}, t)\big]$, where $\mathbb{E}\big[.\big]$ denotes the expectation over all possible feature vectors in the sample space, and $f_S(\mathbf{x}, t)$ denotes some function that measures the discrepancy between the output of S and whether t truly applies to $\mathbf{x}$ (e.g. if $f_S(\mathbf{x}, t)$ is the $0/1-$loss, $\Gamma_S(t)$ is the *true risk* [Sugiyama et al., 2007]). Since we cannot evaluate this expectation (we do not have access to the true distribution of these features), $\Gamma_S(t)$ is not observable, and so it must be inferred from something observable. Standard practice in music autotagging addresses this issue by evaluating S on a test set Φ, and computing an *estimated performance* $\widehat{\Gamma}_S(t)$ (e.g. *empirical risk* in [Sugiyama et al., 2007]). That is, computing a FoM on Φ, and inferring $\Gamma_S(t)$ from this. (Note here again that $\widehat{\Gamma}_S(t)$ is a simplified notation for $\widehat{\Gamma}_{S;\Psi}(t, \Phi)$.) This implicitly assumes that $\widehat{\Gamma}_S(t)$ and $\Gamma_S(t)$ are highly positively correlated.

We define an evaluation to be a *valid indicator of the true performance* $\Gamma_S(t)$ when:

$$[\widehat{\Gamma}_S(t) \text{ good}] \Leftrightarrow [\Gamma_S(t) \text{ high}] \tag{1}$$

and when, for two systems S_1, S_2

$$[\widehat{\Gamma}_{S_1}(t) \text{ better than } \widehat{\Gamma}_{S_2}(t)] \Leftrightarrow [\Gamma_{S_1}(t) \text{ higher than } \Gamma_{S_2}(t)] \tag{2}$$

where $\Leftrightarrow$ is logical equivalence. In other words, (1) says a valid evaluation of S produces a good FoM on t if and only if the true performance of S on t is indeed high; and (2) says a valid evaluation produces a better figure of merit for S_1 than for S_2 on t if and only if the true performance of S_1 is higher than that of S_2 on t.

If, for an evaluation making use of a test set Φ, (1) and (2) do not hold for some tag t, then that evaluation is not a valid indicator of the true performance of S on t. The principal question is no longer, "How good/bad is $\widehat{\Gamma}_S(t)$?", or, "Is $\widehat{\Gamma}_{S_1}(t)$ significantly higher/lower than $\widehat{\Gamma}_{S_2}(t)$?",

2010, Sugiyama et al., 2007].

but now, "Does the evaluation of S in Φ provide a valid indication of its true performance on t?"

3 A method for testing evaluation validity

According to the notion of validity defined in Section 2.4, we now present a method for testing the validity of the evaluation of music autotagging systems. The basic rationale is the following: In experimental conditions where one should expect the true performance of an autotagging system to be relatively stable (see Section 2.3), if its estimated performance varies such that (1) and (2) are violated, then that evaluation is not a valid indicator of the system's true performance.

At its core, our method is based on a systematic search for perceptually indistinguishable test sets, while controlling for the required absence of covariate shift [Sugiyama et al., 2007, Quionero-Candela et al., 2009]. These test sets are obtained by irrelevant transformations of a limited selection of instances in a test set. Our approach is comparable to that of [Szegedy et al., 2014], who test the local generalization capability of their image classification systems. Szegedy et al. show, on three different benchmark datasets (images in their case), that for every test instance that is correctly classified by any of the state-of-the-art systems they studied (deep neural networks), there exists instances in the local vicinity of the original test instance that are perceptually indistinguishable from the original but that are misclassified by the system, in any of the possible classes. They obtain these "adversarial" instances (which they also refer to as "blind spots") by means of "imperceptible" transformations of test instances, found by optimizing the input to maximize the prediction error, while restricting the optimization process to local space around the original test instance. While Szegedy et al. employ a constrained optimization approach to find these adversarial instances, we use a brute force approach to achieve the same results. Furthermore, our aim is not to show the existence of "blind spots", but of testing (1) and (2) for a system.

3.1 Our method

More formally, consider $\mathcal{T} = \{t, \bar{t}\}$, where $\bar{t}$ is the negation of t. For a S, assume $\Gamma_S(t)$ and $\Gamma_S(\bar{t})$ remain constant, i.e., S does not learn about $\mathcal{T}$ after its initial training. Consider a testing dataset Φ of audio signals, each tagged t or $\bar{t}$. Define the transformation of the testing dataset, $\mathcal{F}(\Phi) = \{(F_i(\mathbf{x}_i), \mathcal{X}_i) : i \in \mathcal{I}\}$, where F_i transforms the audio signal $\mathbf{x}_i$, and $\mathcal{I}$ denotes the set of indexes of Φ. Adapting the notion proposed in [Sturm, 2014b], we define $\mathcal{F}(\Phi)$ as an *irrelevant transformation* of Φ if it complies with the following requirements:

- $\forall F_i(\mathbf{x}_i)$, $\mathbf{x}_i$ and $F_i(\mathbf{x}_i)$ are perceptually indistinguishable, i.e., a human describing $\mathbf{x}_i$ as t will also describe $F_i(\mathbf{x}_i)$ as t.
- $\mathcal{F}(\Phi)$ produces no covariate shift with respect to Φ [Sugiyama et al., 2007, Quionero-Candela et al., 2009].

Consider $\widehat{\Gamma}_S(t)$ is significantly better than random. With regards to (1), we thus attempt the following tasks:

A1. Find $\mathcal{F}$ to transform Φ such that $\widehat{\Gamma}_S(t)$ is not significantly better than random.

A2. Find $\mathcal{F}$ to transform Φ such that $\widehat{\Gamma}_S(t)$ is close to perfect.

If we can accomplish A1 and A2, (1) does not hold because $\widehat{\Gamma}_S(t)$ can change between extremes though $\Gamma_S(t)$ stays the same. Procedures A1 and A2 are schematized in figure 1.

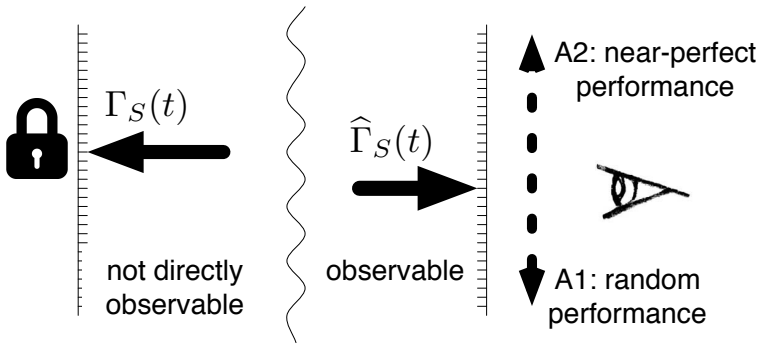


Figure 1: To prove (1) does not hold, while the true performance of S on t, $\Gamma_S(t)$, remains constant (whatever its value), we devise experimental conditions so that its estimator, the figure of merit $\widehat{\Gamma}_S(t)$ takes values ranging from random to close to perfect.

Now, with regards to (2), given two systems S_1 and S_2, we attempt the following:

B1. Find $\mathcal{F}$ to transform Φ such that $\widehat{\Gamma}_{S_1}(t)$ is significantly better than $\widehat{\Gamma}_{S_2}(t)$.

B2. Find $\mathcal{F}$ to transform Φ such that $\widehat{\Gamma}_{S_2}(t)$ is significantly better than $\widehat{\Gamma}_{S_1}(t)$.

If we can accomplish B1 and B2, (2) does not hold because we can make the relative figures of merit of two systems significantly different in either direction while their relative true performance, and ranking, does not change.

3.2 Statistical significance

Task A1 essentially attempts to make the performance of S on Φ decay to the point that it is no longer inconsistent with that of a random system. We thus analyze the behavior of a system that independently picks t for an input with probability p_t (and $\bar{t}$ with probability $1-p_t$). Denote this system by $R(p_t)$. Of the N signals in Φ, consider that there are n_t tagged with t, and $n_{\bar{t}}$ tagged with $\bar{t}$. Let X and Y be random variables for the number of *correct* tags by $R(p_t)$ of t signals and $\bar{t}$ signals, respectively. The probability of $X = x$ is distributed $X \sim Bin(n_t, p_t)$; and of $Y = y$ is distributed $Y \sim Bin(n_{\bar{t}}, 1-p_t)$. The joint probability of $\{X = x, Y = y\}$ is thus:

$$P_{X,Y}(x, y; p_t) = \binom{n_t}{x} p_t^x (1-p_t)^{n_t - x} \binom{n_{\bar{t}}}{y} (1-p_t)^y p_t^{n_{\bar{t}} - y} \tag{3}$$

for $0 \leq x \leq n_t$, $0 \leq y \leq n_{\bar{t}}$, and zero elsewhere.

Now, consider S produces $\{x, y\}$ in Φ. For A1, we test the null hypothesis H_{0A_1}: results at least as good as $\{x, y\}$ are expected from an element of $\{R(p_t) : p_t \in [0, 1]\}$. In other words, observations at least as good as $\{x, y\}$ are consistent with what we expect to be produced by a random system. We test H_{0A_1} by computing:

$$\max_{p_t \in [0,1]} P[X \geq x, Y \geq y; p_t] = \max_{p_t \in [0,1]} \sum_{i=x}^{n_t} \sum_{j=y}^{n_{\bar{t}}} P_{X,Y}(i, j; p_t). \tag{4}$$

and fail to reject H_{0A_1} when this value is greater than the statistical significance parameter α. Recall that our goal with A1 is to show that $\mathcal{F}(\Phi)$ leads to a failure to reject H_{0A_1} though we can reject it for Φ.

For B1 and B2, we must compare the performance of two systems on the same dataset. We count the total number of signals b for which S_1 and S_2 contradict each other, i.e. only one of the systems is wrong. Denote a_{12} the number of signals in the dataset where S_1 makes correct predictions and S_2 is wrong ($b = a_{12} + a_{21}$). If either system is equally likely to be correct (i.e. a_{12} should not be significantly different from a_{21}), then we expect a_{12} to not be significantly different from $b/2$. For B1, the null hypothesis H_{0B_1} is thus $a_{12} = b/2$. Define the random variable $A_{12} \sim Bin(b, 0.5)$ to model a_{12} in b independent trials when S_1 and S_2 are equally likely to be correct when they contradict each other.

Given an observation for a_{12}, we compute the probability that A_{12} is at least as large as a_{12} as:

$$P[A_{12} \geq a_{12}] = \sum_{x=a_{12}}^{b} \binom{b}{x} 0.5^b. \tag{5}$$

If $P[A_{12} \geq a_{12}] < \alpha$, then we reject H_{0B_1}. We follow the same reasoning for B2, and if $P[A_{21} \geq a_{21}] < \alpha$, then we reject H_{0B_2}.

4 Experiments

Here, we first detail our methodology for applying in practice the method defined in Section 3 for evaluating three state-of-the-art systems with three standard datasets. We then present evidence of the irrelevance of the transformations in our experiments. We finally present results on absolute and relative performance of the tested systems, showing that their evaluations are not valid indicators of true performance. In other words, they do not provide valid indicators for concluding whether any of them is objectively good, or better than any other.

4.1 Methodology

We test (1) and (2) for all systems resulting from three state-of-the-art music autotagging approaches crossed with folds of three datasets commonly used for evaluation in music autotagging. We set t as the tag "Vocals", i.e., whether a piece of music includes singing voice or not. We justify this choice by the fact that compared to other possible tags, the tags "Vocals" (t) and "Non-Vocals" ($\bar{t}$) are better defined and more objective relative to other kinds of tags, e.g., genre and emotion, and that it appears in all of our three datasets in some form, e.g., "voice", "gravely voice", or "female singer". This scenario is simpler than the general case of autotagging, but we claim that if the evaluation of a given system can be shown not to provide a valid indication of true performance for such an objective, single-label case, it is not reasonable to assume that the evaluation of that system should be valid in the more subjective and ill-defined general multilabel case (we discuss this further in Section 5). It should also be noted that not only is such a tag suitable to the experimental procedure in this article, but also the actual ability to automatically detect whether a music excerpt includes singing voice or not corresponds to a realistic and very useful problem.

4.1.1 Deflation and inflation procedures

Given a system S and test dataset Φ, we test (1) using what we call "deflation" and "inflation" procedures, that are illustrated in Algorithms 1 and 2 (where $I\mathbf{x} = \mathbf{x}$ is the identity transformation). For deflation, we find irrelevant transformations $\mathcal{F}(\Phi)$ that decrease the number of correct responses by S. As mentioned in Section 3, this is comparable to the procedure of [Szegedy et al., 2014] (in the context of image classification) where for each possible test instance correctly classified by a system they find in its local vicinity an "adversarial" instance that is misclassified, although they are perceptually indistinguishable. In the deflation procedure, we alternate between finding elements of Φ for which S is correct, and transforming these signals in irrelevant ways (as defined in Section 3) to make S respond incorrectly, until the performance of S becomes similar to that of a random

Algorithm 1: Pseudo-code for the deflation procedure.

Initialization: begin

1 $\mathcal{F} \leftarrow \{F_i = I : i \in \mathcal{I}\}$ (Initialize all transformations to identity);

end

repeat

2 $\mathcal{J} \leftarrow \{i \in \mathcal{I} : F_i \in \mathcal{F}\left(S(F_i\mathbf{x}_i) = \mathcal{T}_i\right)\}$ (indices of signals for which S produces correct tags);

3 Produce irrelevant transformation, G;

4 $\mathcal{F} \leftarrow \{F_i = G : i \in \mathcal{J}\} \bigcup \{F_i \in \mathcal{F} : i \in \mathcal{I} \backslash \mathcal{J}\}$ (update set of transformations);

until *the figure of merit of S on the transformed dataset is no better than random*;

Algorithm 2: Pseudo-code for the inflation procedure.

Initialization: begin

1 $\mathcal{F} \leftarrow \{F_i = I : i \in \mathcal{I}\}$ (Initialize all transformations to identity);

end

repeat

2 $\mathcal{J} \leftarrow \{i \in \mathcal{I} : F_i \in \mathcal{F}\left(S(F_i\mathbf{x}_i) \neq \mathcal{T}_i\right)\}$ (indices of signals for which S produces incorrect tags);

3 Produce irrelevant transformation, G;

4 $\mathcal{F} \leftarrow \{F_i = G : i \in \mathcal{J}\} \bigcup \{F_i \in \mathcal{F} : i \in \mathcal{I} \backslash \mathcal{J}\}$ (update set of transformations);

until *the figure of merit of S on the transformed dataset is close to perfect*;

system, according to (4) (with $\alpha = 0.01$). For inflation, we find transformations $\mathcal{F}(\Phi)$ that increase the number of correct responses by S. To do this, we alternate between finding elements of Φ for which S is incorrect, and transforming these signals in irrelevant ways to make S respond correctly. The system's true performance $\Gamma_S(t)$ never changes, but the deflation procedure attempts to make its FoM $\widehat{\Gamma}_S(t)$ worse, while the inflation procedure attempts to make it better. (Note that in both procedures a given signal is transformed at most once and that we seek to transform only a few instances in Φ.) If we are able to produce any FoM of a system just by changing irrelevant aspects of Φ (i.e. transformations do not produce a covariate shift and are perceptually indistinguishable), then (1) does not hold.

We test (2) using the same iterative procedure, but with two systems. Given S_1, S_2 and Φ, we set aside all instances of Φ for which S_1 is correct, but S_2 is not. Then we apply successive transformations to the remaining instances until the performance of S_1 becomes significantly better than that of S_2, according to (5) (with $\alpha = 0.01$). We repeat this procedure, but set aside all instances of Φ for which S_2 is correct and S_1 not, then we apply successive transformations to the remaining instances until the performance of S_2 becomes significantly better than that of S_1.

4.1.2 Signal transformations

Our method in Section 3 does not specify the nature of the irrelevant transformation. This depends on the tag. In our case for Vocals/Non-Vocals tags, examples of transformations that would not be irrelevant are e.g. adding voice to signals without voice, and removing vocals from signals that have voice. Examples of irrelevant transformations for Vocals/Non-Vocals tags may be minor time-stretching and/or pitch-shifting, changes in instrumentation while preserving voice or no voice, minor equalization, and so on. In our experiments here, we use time-invariant filtering, which proceeds as follows. We use the same irrelevant transformation, as well as time-stretching in another work [Sturm et al., 2014]: Specifically, we first build a 96-channel near perfect reconstruction polyphase filterbank.[4] Passing a signal through this filterbank produces 96 signals that when added with unity gain reproduces the original signal with an average reconstruction squared error of -300 dB. We, however, reduce the gains of a randomly selected subset of the 96 channels and then sum the outputs of the filterbank. This subset can be any number of channels, and the attenuation of each channel selected is bounded to be no more than 20 dB. This results in numerous different filters that "equalize" audio signals but preserve the music they embody. Figure 2 shows the magnitude responses of some of these filters. In Section 4.6, we test the irrelevance of these transformations. Audio examples and software code are available on the article's companion webpage (which link is provided in Section 1).

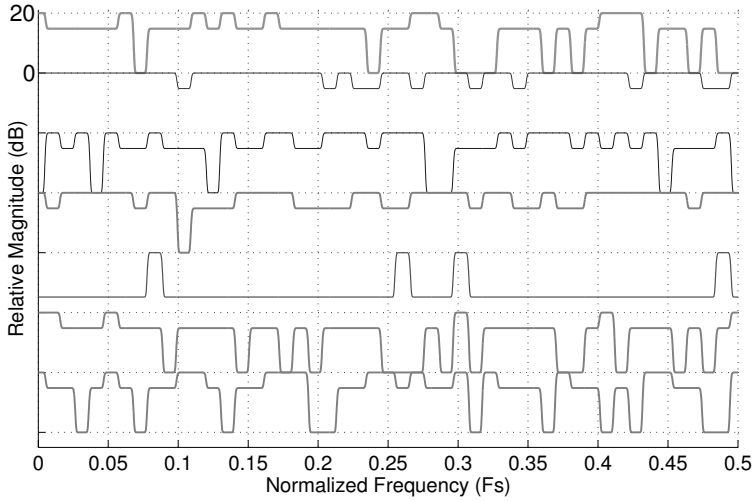


Figure 2: Magnitude responses of a selection of filters used in the deflation procedure. Note that the y-axis is "relative magnitude".

[4] We adopt this code: `http://www.mathworks.com/matlabcentral/fileexchange/15813-near-perfect-reconstruction-polyphase-filterbank`

4.2 Data

We now discuss the data we use, and our preprocessing of it. Table 1 provides data statistics. Data folds are available on the article's companion webpage (link in in Section 1). We use three different datasets, CAL500, a subset of MagnaTagatune, and a subset of the Million Song Dataset, each described below. We reduce the vocabulary of each dataset to the Vocals and Non-Vocals tags, i.e. we keep all instances annotated with a tag corresponding to either Vocals or Non-Vocals tags, we do not consider further the remaining instances. In this process, we favor data *quality* over *coverage*, this has the advantage to make exhaustive listening and checking feasible, offering hence the guarantee of data with no noise in annotations. We correct annotations of the resulting data via a careful listening. The tags Vocals and Non-Vocals are well-defined and relatively objective, mutually exclusive, and always relevant. It is thus straightforward to manually clean and correct annotations of our three datasets with respect to these tags. We split each dataset into folds, and artist filtering [Pampalk et al., 2005,Flexer, 2007] is used to guarantee that no same artist appears in both training and test data.

	CAL500	MagTag5k	MSD24k
Vocals pieces	444	1626	1146
Non-Vocals pieces	58	723	531
Total	502	2349	1677

Table 1: Statistics for the datasets used in experiments.

4.2.1 CAL500

This dataset is a collection of 502 music pieces annotated from a vocabulary of 174 tags. It was first introduced in [Turnbull et al., 2008], it is available online and is widely used in the autotagging literature. When obtained from the original website, we found that all sound files but two were there, although their annotations were. Thus, we corrected this by retrieving the missing songs.

We consider songs originally annotated with tags such as "Female lead vocals", or "Vocals-Gravelly" instances of the Vocals tag (see the full list in Appendix A). There is no explicit Non-Vocals tags in CAL500, so we initially considered all remaining songs as instances of the Non-Vocals tag, and after careful listening, retagged 11 instances from Non-Vocals to Vocals. The dataset is divided in 2 folds.[5]

[5] We chose 2 folds and not 3 (as with the other datasets) because of the relative few Non-Vocals instances (58) in the whole dataset.

4.2.2 MagTag5k

This is a processed version of the original MagnaTagatune dataset (originally 21,642 pieces and a vocabulary of 188 tags [Law et al., 2009]), coping for issues of duplication, synonymy, etc., in the original dataset. Details about the preprocessing applied on that dataset can be found in [Marques et al., 2011]. This dataset consists of 5,259 music pieces annotated from a vocabulary of 137 tags, and is available online.[6]

We assign the Vocals tag to songs annotated with the tags "female.singing", "male.singing", or "singing". We assign the Non-Vocals tag to songs annotated with the tags "no.singing". This yields 2,393 songs, which we check by careful listening, after which the final dataset contains 2,349 instances, see Table 1. The dataset is divided in 3 folds.

4.2.3 MSD24k

We designed the MSD24k dataset for in-house experiments in music autotagging, with the main objective to set up a dataset, comprising the audio data, with tags of relatively good quality and with the highest density of annotations possible (i.e. imposing a lower limit on the number of tags per music piece). As this article is the first publication referring to it, we now describe the procedure followed in its creation.

This dataset is based upon the subset of the Million Song Dataset (MSD) [Bertin-Mahieux et al., 2011] for which the MSD website provides[7] `Last.fm` tags associated to its tracks (943,347 tracks). In order to cope with the significant problem of noise in `Last.fm` tags [Lamere, 2008], we follow the same rationale as [Tingle et al., 2010] and focus on tags with clear musical meaning, as defined by teams of musicologists of the Music Genome Project at the Pandora Internet radio. We therefore generate a *relevant tag vocabulary* $\mathcal{T}$ consisting of the overlap between Pandora tags (gathered from the CAL10k dataset [Tingle et al., 2010]) and existing `Last.fm` tags from MSD. This vocabulary contains 708 tags. Retrieving the music pieces from MSD with at least 1 tag in $\mathcal{T}$ yields a total of 257,387 pieces. We then keep only pieces with *at least 4 tags per piece*, lowering the total number of pieces to 60,769. Of these, we were only able to retrieve 30 s snippets of 36,000 pieces in mp3 format. Removing duplicates yields 26,277 pieces. We finally remove the pieces corresponding to the "list of MSD {song ID, track ID} pairs that should not be trusted" (list available online).[8] This yields a final amount of 23,740 music pieces annotated from a vocabulary of 265 tags.

We assign the Vocals tag to songs annotated with tags such as "A breathy male lead vocalist", or "A distinctive male lead vocalist". Appendix A lists

[6] `http://tl.di.fc.ul.pt/t/magtag5k.zip`

[7] on `http://labrosa.ee.columbia.edu/millionsong/lastfm`

[8] `http://labrosa.ee.columbia.edu/millionsong/sites/default/files/tasteprofile/sid_mismatches.txt`

the full tag list. As for the CAL500 dataset, there is no explicit Non-Vocals tags in MSD24k, however in that case the dataset size makes very difficult an exhaustive listening. Therefore, we recur to the following heuristics to select Non-Vocals instances. We divide the dataset in 2 groups: Group A made up of songs in the Vocals tag, and Group B made up of the remainder. We then rank all tags according to their representativeness of both groups, from "occuring mostly in songs from Group A", to "occuring mostly in songs from Group B". We then take a random sample of 1000 songs annotated only with the most representative tags of Group B. After careful listening to these songs, we keep 531 instances of the Non-Vocals tag. (Note here that with this procedure, we favor quality over coverage of Non-Vocals instances.) The dataset is divided in 3 folds.

4.3 Building Music Autotagging Systems

We use three different approaches to build music autotagging systems. The first, SVMBFFs, combines bags of frames of features (BFFs) and a support vector machine classifier (SVM). The second, VQMM, first codes a signal using vector quantization (VQ) in a learned codebook, and then estimates conditional probabilities in first-order Markov models (MM). The third, SRCAM, employs sparse representation classification to approximate a high-dimensional psychoacoustically-motivated frequency modulation feature. Below, we discuss each approach in more detail.

4.3.1 SVMBFFs

This approach, a variant of one proposed by [Ness et al., 2009], trains a linear SVM to output probabilities from an input BFFs, from which tags are selected. The BFFs, which are 68-dimensional vectors, are means and standard deviations computed from texture windows of 30 s of analysis frames of 23.2 ms duration (and overlapped by 50%). The 17 low-level features extracted from each frame are: zero crossing rate, spectral centroid, roll-off and flux, and the first 13 mel-frequency cepstral coefficients (MFCCs). SVMBFFs trains an SVM by a "normalized" training dataset of BFFs, i.e., where each dimension of the set of transformed BFFs lies in $[0, 1]$. We use the SVMBFFs implementation available in the MARSYAS framework.[9]

4.3.2 VQMM

This approach computes the 13 MFCCs after the zeroth with an analysis frame of 93 ms using the YAAFE toolbox.[10] Analysis frames are overlapped by 50%. Given the feature vectors $\{\mathbf{f}_1, \mathbf{f}_2 \ldots, \mathbf{f}_n\}$ extracted from an input

[9] MARSYAS can be downloaded here: `http://marsyas.info/`. We use default settings of `bextract` and `kea` v.5099.

[10] `http://yaafe.sourceforge.net/`

signal, VQMM first expresses it as an ordered code $\{w_1, w_2, \ldots, w_n\}$ in a codebook $\mathcal{C}$, then computes a probability of observing this code in each of a set of duples of models $\{(M_t, \overline{M}_t) : t \in \mathcal{T}\}$, and finally selects a set of tags from $\mathcal{T}$ based on maximum likelihood. The duple of models $(M_t, \overline{M}_t)$ is composed of a model M_t trained on coded features for which the tag $t \in \mathcal{T}$ is relevant, and a model $\overline{M}_t$ trained on coded features for which it is not relevant. In our case, M_t models "Vocals", and $\overline{M}_t$ models "Non-Vocals". VQMM computes the probability of observing the ordered code $\{w_1, w_2, \ldots, w_n\}$ in the model of tag $t \in \mathcal{T}$, $P_{M_t}(w_1, w_2, \ldots, w_n)$, as well as its complement, $P_{\overline{M}_t}(w_1, w_2, \ldots, w_n)$. If $P_{M_t}(w_1, w_2, \ldots, w_n) > P_{\overline{M}_t}(w_1, w_2, \ldots, w_n)$, VQMM selects t as a tag for the input.

VQMM builds a codebook by first grouping all features extracted from the signals in a training dataset into $K = 75$ clusters using k-means [Gersho and Gray, 1991] –though other unsupervised approaches could be used– and then pairing the K centroids of the clusters with codewords. To code a feature vector in terms of the codebook, VQMM selects the codeword of the nearest (in a Euclidean sense) centroid in the codebook.

VQMM builds a model under the assumption that the ordered code is a first-order Markov process, i.e., all pairs of elements from an ordered code $\{w_1, w_2, \ldots, w_n\}$, except for those that are subsequent, are independent. The log joint probability of this code in M_t thus becomes

$$\log P_{M_t}(w_1, w_2, \ldots, w_n) = \log P_{M_t}(w_1) + \sum_{i=1}^{n-1} \log P_{M_t}(w_{i+1}|w_i). \qquad (6)$$

VQMM trains M_t by estimating the set of conditional probabilities $\{P_{M_t}(w_i|w_j) : w_i, w_j \in \mathcal{C}\}$, as well as $\{P_{M_t}(w_i) : w_i \in \mathcal{C}\}$, from coded feature vectors extracted from the training instances for which t is a relevant tag. VQMM uses the coded features of all other signals to train $\overline{M}_t$. More details can be found in [Langlois and Marques, 2009].[11]

4.3.3 SRCAM

This approach, a variant of one proposed by [Panagakis et al., 2009, Sturm and Noorzad, 2012] and [Sturm, 2012], uses sparse representation classification (SRC) [Wright et al., 2009] of auditory temporal modulation features (AM). Here, we extend it to a multilabel classifier. Given the *dictionary* of feature atom-tag atom duples $\{(\mathbf{d}_i, \mathbf{t}_i/\|\mathbf{t}_i\|_2) : i \in \mathcal{I}\}$, SRCAM approximates a feature vector $\mathbf{f}$ as a linear combination of a small number of feature atoms, and then produces a tag vector $\mathbf{t}$ by thresholding a linear combination of the tag atoms.

[11] Source code is available at `https://bitbucket.org/ThibaultLanglois/vqmm`.

More formally, SRCAM first solves

$$\min_{\mathbf{s}} \|\mathbf{s}\|_1 \text{ subject to } \left\| \frac{\mathbf{f}}{\|\mathbf{f}\|_2} - [\mathbf{d}_1|\mathbf{d}_2|\cdots]\mathbf{s} \right\|_2^2 \leq \epsilon^2 \tag{7}$$

then uses the solution $\mathbf{s}$ to produce the linear combination of tag atoms $\mathbf{w} = [\mathbf{t}_1/\|\mathbf{t}_1\|_2 \,|\, \mathbf{t}_2/\|\mathbf{t}_2\|_2 \,|\, \cdots]\mathbf{s}$, and finally produces from this the tag vector $\mathbf{t} = T_\lambda(\mathbf{w}/\|\mathbf{w}\|_\infty)$, where $T_\lambda(\cdot)$ is a threshold operator, its ith element defined

$$[T_\lambda(\mathbf{w}/\|\mathbf{w}\|_\infty)]_i = \begin{cases} 1, & [\mathbf{w}]_i/\|\mathbf{w}\|_\infty > \lambda \\ 0, & \text{else.} \end{cases} \tag{8}$$

The non-zero dimensions of $\mathbf{t}$ correspond to the tags in $\mathcal{V}$ considered relevant for annotating the input signal.

SRCAM defines the dictionary from a training feature-tag vector dataset by first constructing a matrix of the features, $\mathbf{F} = [\mathbf{f}_1|\mathbf{f}_2|\ldots]$, finding the maximum and minimum of each dimension, defined as column vectors $\max \mathbf{F}$ and $\min \mathbf{F}$, respectively, and then computing the matrix of *normalized* feature atoms

$$\mathbf{D} = [\mathbf{d}_1|\mathbf{d}_2|\cdots] = \text{diag}(\max \mathbf{F} - \min \mathbf{F})(\mathbf{F} - \mathbf{1}[\min \mathbf{F}]^T). \tag{9}$$

Normalization guarantees that each dimension of $\mathbf{D}$ is in $[0, 1]$.

The particulars of our implementation of SRCAM are as follows. We solve (7) using SPGL1 [van den Berg and Friedlander, 2008], and define $\epsilon^2 = 0.01$ and 200 iterations from experimentation. For thresholding (8), we define $\lambda = 0.25$ from experimentation. We compute features from contiguous segments of about 27.7 s duration in a signal. Specifics about computing AMs are given in [Sturm, 2012].

4.3.4 Baseline Results

We test these systems on the CAL500 dataset, but restricted to the 97 most frequent tags (as done in [Miotto et al., 2010, Xie et al., 2011, Nam et al., 2012, Coviello et al., 2012]). We use 5-fold cross-validation, and compute (as is standard in autotagging research) the mean per-tag precision, recall and F-score of all systems. Table 2 shows good FoM of our three systems, which are on-par with those of four other state-of-the-art approaches (included in the table). We also test all systems on the three datasets, restricted to the tag vocabulary of Vocals and Non-Vocals. Table 3 shows very good results for these systems.

4.4 On absolute performance (tasks A1 and A2 in practice)

We now perform tasks A1 and A2 using the methodology in Section 4.1. For a given system S (which is already trained on a subset of data folds) and a

	CAL500 (97 tags)		
	P	R	F
SVMBFFs	0.40	0.40	0.40
VQMM	0.38	0.46	0.42
SRCAM	0.34	0.57	0.42
HEM-DTM [Coviello et al., 2012]	0.45	0.22	0.26
[Miotto et al., 2010]	0.44	0.23	0.30
[Xie et al., 2011]	0.45	0.23	0.30
[Nam et al., 2012]	0.48	0.26	0.29

Table 2: Average per-tag precision, recall and F-score of the three systems, compared to recent systems, on CAL500 restricted to the 97 most frequent tags, 5-fold cross-validation procedure.

		CAL500			MagTag5k			MSD24k		
		P	R	F	P	R	F	P	R	F
S_1	V	0.92 ± 0.02	0.99 ± 0.00	0.95 ± 0.01	0.88 ± 0.01	0.91 ± 0.02	0.89 ± 0.01	0.89 ± 0.01	0.92 ± 0.01	0.91 ± 0.01
	NV	0.78 ± 0.04	0.33 ± 0.17	0.45 ± 0.18	0.79 ± 0.03	0.72 ± 0.02	0.75 ± 0.01	0.82 ± 0.02	0.77 ± 0.03	0.80 ± 0.02
S_2	V	0.93 ± 0.01	0.96 ± 0.02	0.95 ± 0.01	0.85 ± 0.02	0.85 ± 0.03	0.85 ± 0.01	0.85 ± 0.01	0.80 ± 0.00	0.83 ± 0.01
	NV	0.63 ± 0.11	0.48 ± 0.09	0.54 ± 0.02	0.66 ± 0.02	0.67 ± 0.06	0.66 ± 0.02	0.62 ± 0.01	0.71 ± 0.03	0.66 ± 0.02
S_3	V	0.94 ± 0.01	0.95 ± 0.02	0.95 ± 0.01	0.88 ± 0.01	0.92 ± 0.02	0.90 ± 0.004	0.89 ± 0.01	0.94 ± 0.01	0.91 ± 0.004
	NV	0.60 ± 0.12	0.55 ± 0.05	0.57 ± 0.08	0.80 ± 0.03	0.73 ± 0.04	0.76 ± 0.01	0.86 ± 0.01	0.74 ± 0.03	0.80 ± 0.02

Table 3: Average $\pm$ standard deviation, for Precision, Recall and F-Score for the 3 systems on CAL500, MagTag5k and MSD24k (respectively with 2-fold, 3-fold and 3-fold cross-validations). Vocabulary restricted to Vocals ("V" rows) and Non-Vocals ("NV" rows) . S_1 is SVMBFFs, S_2 is VQMM, and S_3 is SRCAM.

test dataset Φ (remaining fold of dataset), we aim to find the set of irrelevant transformations $\mathcal{F}_{def}(\Phi)$ (for "deflation") and $\mathcal{F}_{inf}(\Phi)$ (for "inflation") such that S performs no better than random for $\mathcal{F}_{def}(\Phi)$, and S performs close to perfectly for $\mathcal{F}_{inf}(\Phi)$. Section 4.6 below confirms the irrelevance of our transformations using covariate shit and listening tests.

Figure 3 shows the FoM of three SVMBFFs systems, trained on three combinations of two MSD24k folds and tested on the three respectively remaining folds. FoM is plotted versus iterations of the deflation and inflation procedures applied to the test set. On all three folds, we see that our procedures yield clear decrease and increase in FoM in very few iterations.

Figure 4 shows the FoM of three SRCAM systems trained on one CAL500 fold (black line), two MagTag5k folds (blue line) and two MSD24k folds

		CAL500		MagTag5k			MSD24k		
		Fold 1	Fold 2	Fold 1	Fold 2	Fold 3	Fold 1	Fold 2	Fold 3
S_1	$\mathcal{F}_{def}(\Phi)$	√	√	√	√	√	√	√	√
	$\mathcal{F}_{inf}(\Phi)$	1.0	1.0	0.89	0.96	0.89	0.99	0.93	0.93
S_2	$\mathcal{F}_{def}(\Phi)$	√	√	√	√	√	√	√	√
	$\mathcal{F}_{inf}(\Phi)$	**0.95**	0.97	**0.97**	0.97	0.98	**0.96**	0.95	0.95
S_3	$\mathcal{F}_{def}(\Phi)$	√	√	√	√	√	√	√	√
	$\mathcal{F}_{inf}(\Phi)$	0.98	0.97	0.98	0.99	0.99	0.99	0.99	0.99

Table 4: Effect of the deflation and inflation procedures applied to test sets. S_1 is SVMBFFs, S_2 is VQMM, and S_3 is SRCAM. Columns correspond to the test folds (corresponding training data are the remaining folds). √ denotes cases where a system with initial performance superior to random ($p < \alpha = 0.01$ in (4)) performs consistently to random after deflation of the test set. Reported average per-tag F-scores after inflation of the test sets ($\mathcal{F}_{inf}(\Phi)$ rows) are close to perfect. In bold, results obtained with data which train/test divergence is reported in the second column of Table 5.

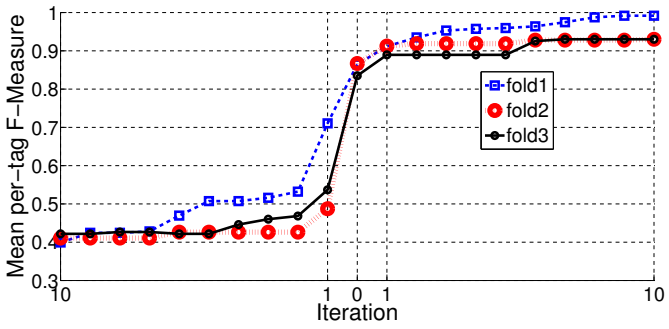


Figure 3: Mean per-tag F-measure (average over Vocals and Non-Vocals) with respect to ten successive iterations of the deflation procedure (iterations left to the origin) and inflation procedure (iterations right to the origin), as detailed in Section 4.1, for three SVMBFFs systems tested on three different folds of MSD24k. F-measure at iteration 0 for the three folds (≈ 0.85) corresponds to average performance of SVMBFFs on MSD24k as can be seen on Table 3.

(red line) respectively, and tested on the remaining fold of the respective dataset. The line corresponding to each system represents change in FoM with respect to successive transformations of the test set. In other words, the opposite ends of a given line correspond to the FoM obtained either after deflation or inflation of the test set. One can see that the performance of all systems can take on drastically different values after few iterations of irrelevant transformations. Namely, the performance of each system can be significantly better than random (outside region demarcated by black lines), to no better than random (inside region, according to (4)).

Table 4 reports results for all systems using SVMBFFs, VQMM and

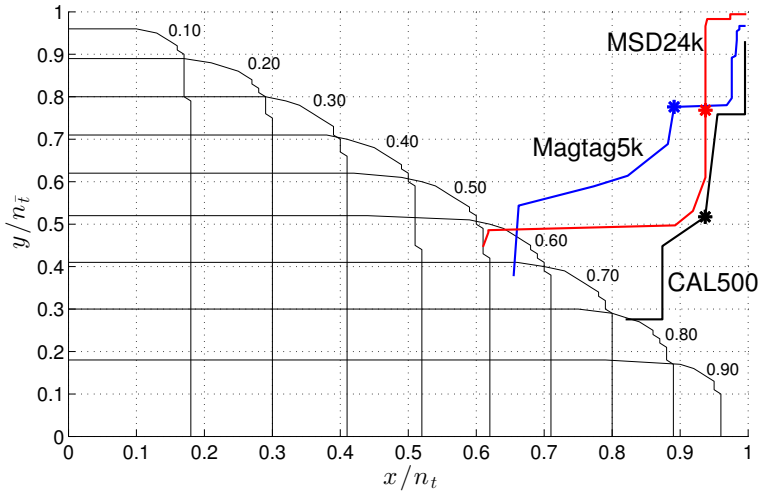


Figure 4: For three systems created using the SRCAM approach, we are able to transform the test data –CAL500 (black), MagTag5k (blue), and MSD24k (red)– such that their performance is near perfect ($\mathcal{F}_{inf}(\Phi)$, top right corner), or consistent with that expected from a random system $R(p_t)$ ($\mathcal{F}_{def}(\Phi)$, within thin black lines, where $p > \alpha = 0.01$) that randomly picks t with probability p_t (illustrated here between 0.10 and 0.90, in steps of 0.10) and $\bar{t}$ with probability $1 - p_t$. Each star marks the "starting position" of the system. x/n_t is the ratio of correctly classified instances of Vocals, $y/n_{\bar{t}}$ is the ratio of correctly classified instances of Non-Vocals.

SRCAM approaches, on all folds of the three datasets. Each cell in the table corresponds to a system built using one of the three approaches, trained on some data folds of a given dataset, and tested on the remaining fold. Results correspond to either the deflation or inflation procedures. The performance of each system can vary between almost perfect to no better than random, while the diversity of experimental conditions has no effect on whether a given piece of music includes singing voice or not, and is perceived as such.

4.5 On relative performance (tasks B1 and B2 in practice)

We now perform tasks B1 and B2 using the methodology in Section 4.1. For two given systems S_i and S_j (already trained on a subset of data folds) and a test dataset Φ (remaining fold), we aim to find a transformation $\mathcal{F}_i$ such that S_i performs significantly better (according to (5)) than S_j on $\mathcal{F}_i(\Phi)$, and another transformation $\mathcal{F}_j$ such that the opposite is true on $\mathcal{F}_j(\Phi)$.

After conducting experiments for all possible pairwise comparisons of any two systems among SVMBFFs, VQMM, and SRCAM, on any possible test set among each of the three datasets we use, we can report that it is always possible, in a few iterations, to find an irrelevant transformation of

any test set so that any two systems are alternatively the best.[12]

4.6 Testing the irrelevance of the transformations

4.6.1 On the irrelevance of the transformations with respect to covariate shift

In our experimental procedure, measuring covariate shift is important for verifying irrelevance of the transformations. We need to make sure that there is no significant divergence between the feature distributions of train and test data. For this, we follow the method proposed by [Ben-David et al., 2010]. They show that an upper bound on the divergence $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}')$ between two distributions $\mathcal{D}$ and $\mathcal{D}'$ can be estimated from an empirical divergence $\hat{d}_{\mathcal{H}}(\mathcal{U},\mathcal{U}')$ computed from finite samples $\mathcal{U}$ and $\mathcal{U}'$ of these distributions.

The method for computing $\hat{d}_{\mathcal{H}}(\mathcal{U},\mathcal{U}')$ consists in labelling each instance $x \in \mathcal{U}$ with 0, and each instance $x \in \mathcal{U}'$ with 1. Then training classifiers[13] $h \in \mathcal{H}$ to discriminate between instances of $\mathcal{U}$ and $\mathcal{U}'$. In a testing phase, one can then compute a confusion matrix for each classifier h and compute $\hat{d}_{\mathcal{H}}(\mathcal{U},\mathcal{U}')$ as follows (lemma 2 in [Ben-David et al., 2010]):

$$\hat{d}_{\mathcal{H}}(\mathcal{U},\mathcal{U}') = 2\left(1 - \min_{h\in\mathcal{H}}\left[\frac{1}{m}\sum_{x:h(x)=0} I[x \in \mathcal{U}] + \frac{1}{m}\sum_{x:h(x)=1} I[x \in \mathcal{U}']\right]\right) \quad (10)$$

where m is the number of instances in $\mathcal{U}$ and $\mathcal{U}'$ and $I[x]$ indicates class membership of x (i.e. $I[x \in \mathcal{U}] = 1$ if $x \in \mathcal{U}$). Smaller values in (10) refer to smaller divergence. As noted in [Ben-David et al., 2010], it is not feasible to compute (10) with the minimum over *all possible* classifiers $h \in \mathcal{H}$. In our experiments below, we therefore compute the minimum over ten different classifiers (which we choose to be linear perceptrons).

An upper bound on $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}')$ is then given by the following equation (lemma 1 in [Ben-David et al., 2010]):

$$d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}') \leq \hat{d}_{\mathcal{H}}(\mathcal{U},\mathcal{U}') + 4\sqrt{\frac{d\log(2m) + \log(2/\delta)}{m}} \quad (11)$$

where d is $\mathcal{H}$'s VC dimension [Ben-David et al., 2010], and $\delta \in (0, 1)$ is a confidence parameter.

In the case where the samples $\mathcal{U}$ and $\mathcal{U}'$ are drawn from the same distribution, for instance if $\mathcal{U}$ is a sample of a training fold Ψ and $\mathcal{U}'$ a sample of a test fold Φ of the same dataset, the classifiers h should do a bad job a

[12] See the article's companion webpage (link in in Section 1) for results and their reprocuction (i.e. 3 systems $*$ 2 conditions $*$ (2+3+3) folds = 48 comparisons in total).

[13] $\mathcal{H}$ is a class of functions from features to tag, which, for consistency with the rest of this article, we refer to as a set of classifiers (e.g. linear perceptrons). The correct naming would be a "hypothesis class" [Ben-David et al., 2010].

discriminating between instances of $\mathcal{U}$ and $\mathcal{U}'$. $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}')$ should therefore be low. In our experiments below, we precisely compare the divergence in such cases (namely when no data is transformed) to the divergence when some data is transformed by inflation or deflation.

The first column of Table 5 corresponds to cases where we define $\mathcal{U}$ as 100k randomly selected frames from one data fold of a given dataset, and $\mathcal{U}'$ as 100k randomly selected frames of the complementing fold(s) of that dataset.[14] We then use half of $\mathcal{U}$ and half of $\mathcal{U}'$ for training simple linear perceptrons, and the remaining halves for computing (10). Two trials were done for each dataset. In these cases, in the first column of Table 5, in each line, the data is coming from a single dataset, and *no* instance is transformed, the divergence values obtained are therefore representative of standard cases of autotagging evaluation (i.e. cross-validation) where one can consider that there is no significant divergence in feature distributions of train and test data, i.e. no covariate shift. The inter-row differences provide examples of non-significant variability in the computation of the divergence.[15]

The second column of Table 5 corresponds to cases where we define $\mathcal{U}'$ as 100k randomly selected frames of the *transformed* fold of a given dataset (namely the transformed fold used for test in inflation and deflation experiments which results are reported in bold in Table 4), and where we define $\mathcal{U}$ as 100k randomly selected frames from the complementing data fold(s) of that dataset. The second column shows that when applying transformations (either inflation or deflation) to the test set, the upper bounds for the divergence between training and test sets are relatively low, and sensibly the same as when no transformation is applied (i.e., in the first column). This provides evidence of the irrelevance of the transformations with respect to covariate shift.

		Ψ *vs* Φ		Ψ *vs* $\mathcal{F}(\Phi)$
CAL500	trial 1	0.34	$\mathcal{F}_{inf}(\Phi)$	0.35
	trial 2	0.38	$\mathcal{F}_{def}(\Phi)$	0.39
MagTag5k	trial 1	0.40	$\mathcal{F}_{inf}(\Phi)$	0.34
	trial 2	0.37	$\mathcal{F}_{def}(\Phi)$	0.36
MSD24k	trial 1	0.24	$\mathcal{F}_{inf}(\Phi)$	0.26
	trial 2	0.27	$\mathcal{F}_{def}(\Phi)$	0.39

Table 5: Upper bounds for $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}')$, computed as (11). $\mathcal{F}_{inf}(\Phi)$ and $\mathcal{F}_{def}(\Phi)$ rows correspond to inflation or deflation procedures applied to the test set which corresponding performances are reported in bold in Table 4.

[14] Recall that for computing (10), the labelling of instances $x \in \mathcal{U}$ with 0 and $x \in \mathcal{U}'$ with 1 have nothing to do with Vocals and Non-Vocals tags. $\mathcal{U}$ and $\mathcal{U}'$ are random frames from Vocals and Non-Vocals instances.

[15] Divergence upper bounds are $\neq 0$ because of the second term in the right-hand side of (11) and by the fact that a linear perceptron is a weak classifier. A better classifier would probably give tighter bounds.

4.6.2 On the perceptual irrelevance of the transformations

A key aspect in our experiments relies on our assumption of perceptual irrelevance of the deflation and inflation procedures. In order to verify this assumption, we perform a listening test, where 152 subjects are asked to rate 32 audio stimuli with respect to whether they contain singing voice or not. Stimuli are representative of those used in experiments with autotagging systems in Sections 4.4 and 4.5, i.e. half of the stimuli are "originals", while the other half are transformed according to deflation or inflation procedures. Results show that recognition of singing voice is very good, i.e. $\approx 98\%$, and that there is no significant effect of the condition (original or transformed). More details are available in Appendix B.

5 Summary and Discussion

In this article, we tackle the issue of validity in the evaluation of music autotagging systems. For a given music autotagging system, a valid evaluation means that there is a high positive correlation between its figure of merit and its true performance on the task for which it has been designed. This is essential for making relevant conclusions about a system's performance in laboratory conditions (and all the more in real-world conditions). Validity is, more generally, paramount to guarantee continued improvements in autotagging system research and development. Our main contributions in this paper are the formalization of the notion of validity in autotagging evaluation and the proposal of a method for testing it (with available code), which centers on the control of experimental conditions via irrelevant transformations of input signals.

We demonstrate the use of our method with three autotagging systems in a simple two-class setting (i.e. recognizing the presence or absence of singing voice in an excerpt). We find we can make all three perform as well or as poorly as we like by irrelevant transformations. Although these systems initially appear to be on-par with current state-of-the-art, their FoM do not provide valid indicators of their true performance on the task of recognizing the presence or absence of singing voice in an excerpt, and do not provide valid indicators for comparing them in that task.

An important point to clarify is that our method does not aim to answer questions regarding system performance in the real world. It is designed first and foremost to answer questions about what the systems have learned to do. And our conclusions are limited to particular datasets. In other words, our experiments aim to answer whether the observation of the systems' FoM, or comparisons thereof, warrant any conclusion about the actual capacity of these systems to annotate CAL500, MagTag5k, or MSD24k data with the concept of singing voice. We claim that our experiments provide evidence that this is in fact not the case. Questioning whether these systems would

be able to apply that concept in the real world (where e.g. covariate shift would probably happen) is another question altogether, which we do not address in this article.

Since we consider a special case of autotagging that is simpler than the general case of multi-label classification, i.e., we consider only music labeled using two mutually exclusive tags, "Vocals" and "Non-Vocals", the generality of our work here may appear limited; the autotagging systems used in this article are indeed not designed only for this two-class problem, but for multi-label classification (including these two classes nevertheless). We also do not claim that the evaluation of these systems is necessarily uninformative for any possible tag. Instead, we just show that even for what should be a simple case for these systems, it is not possible to conclude upon the degree to which they have learned to perform the task. We do claim that this sheds doubt on knowledge we could obtain with certitude in more difficult cases. For instance, if we cannot make valid conclusions about these systems' ability to recognize singing voice, how could these evaluation approaches suddenly serve for solid conclusions on the finer, and more subjective tags like "Vocals-Aggressive," "Vocals-Call & Response," "Vocals-Falsetto," and "Vocals-Rapping"?

It is important to clarify that, although our method uses signal transformations at its core, it is fundamentally different from robustness testing. We ask a different scientific question. While robustness testing asks "How does the performance of S change in condition X?", we ask "Does the evaluation of S provide a valid indication of its true performance?" More than testing the robustness of a particular autotagging *system*, our claims in this article are relative to the validity of the *evaluation* itself. In other words, we use a similar machinery as robustness tests, but only as part of a method whose aim is to test evaluation validity. Further, in existing work on robustness testing [Sigurdsson et al., 2006, Jensen et al., 2009, Urbano et al., 2014, Gouyon et al., 2006, Mauch and Ewert, 2013], experimental conditions are made increasingly more challenging, and decreasing performance is assumed to illustrate disruptibility of a system and its inability to complete its task *under all possible conditions*. Robustness testing is thought to highlight e.g. possibly overestimated FoM, but representative FoM nevertheless. Thus the comparison and ranking of several systems is still thought to be possible and informative. In contrast, we claim that the diverse experimental conditions (i.e. all possible $\mathcal{F}(\Phi)$, including no transformation at all) should not reflect significantly on the behavior of systems if they are pairing audio signals with tags in a meaningful way. Under these experimental conditions, we showed that not only the estimated performances of three systems can drop to random, but it can also ascend to almost perfect, thus providing no valid indication of true performance of these systems on a simple task, and hence uninformative with regards to these systems' ranking.

The erratic behavior of systems' FoM under our experimental conditions

does not mean that the performance measure itself (e.g. the average per-tag F-score) is to blame, or that the systems we consider are unable to learn from data. Instead, it may indicate that what the systems are learning may not necessarily be what they are assumed to have learnt, i.e. the particular dimensions of interest to the evaluator (e.g. the presence or absence of singing voice). Observing correlations between some characteristics of music audio signals and a particular tag cannot by itself lead to the conclusion that the former are necessarily relevant to the latter. Such correlations are just an indication that the former *may* be relevant to the latter [Aldrich, 1995]. In other words, irrelevant characteristics may be confounded with the dimensions of interest [Sturm, 2014b]. Indeed it is likely that the autotagging systems we consider are able to learn from training data an uncontrolled (and unidentified) confounding variable, rather than the presence or absence of singing voice. This factor is highly correlated with the presence/absence of singing voice on the datasets we considered, hence explaining the good FoM in Table 3. (Note that a similar argument on the impact of confounding variables on estimated performance was made in previous MIR work, in the particular case of artist and album effects [Pampalk et al., 2005, Flexer, 2007].) Although our transformations are irrelevant to singing voice, they do affect that confounding variable, hence explaining the large variations in FoM we see e.g. in Table 4. If, for instance, all excerpts tagged "Vocals" in a dataset are loud, and all excerpts tagged "Non-Vocals" are quiet, then the evaluation of a system exploiting only loudness to discriminate between the two will measure the system to be perfect, yet providing no validity for drawing reasonable conclusions on the true performance of that system for actually recognizing singing voice in that dataset.

How could one reliably conclude anything about the ability of a given autotagging system to perform the task at hand? Before being a question of which statistical test to use, or which figures of merit to avoid, it is first and foremost a matter of the design, implementation, and analysis of an evaluation that is valid with respect to estimating true performance. An evaluation is either valid or invalid with respect to the question one is attempting to address –no matter the actual results of the evaluation. [Urbano et al., 2013] discuss several important notions of validity in scientific experiments, and how they relate to MIR. Another critical component is formalizing evaluation [Bailey, 2008,Sturm, 2014a]. In this paper we build on previous research by proposing a method (and code) for testing validity in music autotagging experiments, adapting the method in [Sturm, 2014b], which is reproduced independently in [Szegedy et al., 2014] for image tagging.

Another important point to reiterate here is that what is general in our proposed method for evaluation validity is the notion of "irrelevant transformation," not the particular transformation itself (i.e. our time-invariant filtering). Indeed, the irrelevance of a particular transformation largely depends on the task at hand. In this article, for the purpose of demonstrating

the use of our method, we show that our time-invariant filtering is irrelevant to the specific task of Vocals/Non-Vocals autotagging. Time-stretching, e.g., may have been another option for that task [Sturm et al., 2014]. On the other hand, time-invariant filtering would probably not be appropriate to our method if the task at hand were to annotate music audio signals with tags related e.g. to audio production quality, such as "low-fi" vs. "hi-fi" for instance. In other words, future extensions of the work presenting here may call for different transformations.

Future work will look into which other irrelevant transformations can be designed for testing the validity of evaluation in other MIR tasks. We believe that building our method into MIREX-like campaigns would also be of interest. [Bailey, 2008] provides a very interesting starting point to further work on the formalization of the notion of confounds in MIR research. Another interesting avenue for future work is the adaptation to music autotagging of existing research on the design of systems that can be used in different conditions than those under which they were developed. For instance, an adaptation of our method may be used to attempt to train better systems, as suggested in [Szegedy et al., 2014]. Namely, one could train systems on datasets "enriched" by carefully designed perturbations of instances. Other methods to train systems able to cope with different conditions than those under which they were developed may be adapted from [Quionero-Candela et al., 2009, Pan and Yang, 2010, Ben-David et al., 2010, Sugiyama et al., 2007].

Acknowledgments

FG (fgouyon@inesctec.pt) and NH are with INESC TEC, Porto, Portugal. BLS is with the Audio Analysis Lab, AD:MT, Aalborg University Copenhagen, Denmark. JLO is with INESC TEC and FEUP, Porto, Portugal. TL is with the Science Faculty of Lisbon University, Portugal. FG acknowledges the support of the Media Arts and Technologies project (MAT), NORTE-07-0124-FEDER-000061, financed by the North Portugal Regional Operational Programme (ON.2 O Novo Norte), under the National Strategic Reference Framework (NSRF), through the European Regional Development Fund (ERDF), and by national funds through the Portuguese funding agency Fundação para a Ciência e a Tecnologia (FCT). BLS acknowledges the support of Independent Postdoc Grant 11-105218 from Det Frie Forskningsråd. We thank Matthew Davies, Marcelo Caetano, João Gama, Guy Madison, Doug Turnbull and anonymous reviewers for useful discussions.

References

[Aldrich, 1995] Aldrich, J. (1995). Correlations Genuine and Spurious in Pearson and Yule. *Statistical Science*, 10(4):364–376.

[Bailey, 2008] Bailey, R. A. (2008). *Design of comparative experiments.* Cambridge: Cambridge University Press.

[Ben-David et al., 2010] Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. (2010). A theory of learning from different domains. *Mach. Learn.*, 79(1-2):151–175.

[Berenzweig et al., 2004] Berenzweig, A., Logan, B., Ellis, D. P. W., and Whitman, B. (2004). A large-scale evaluation of acoustic and subjective music-similarity measures. *Computer Music J.*, 28(2):63–76.

[Bertin-Mahieux et al., 2008] Bertin-Mahieux, T., Eck, D., Maillet, F., and Lamere, P. (2008). Autotagger: A model for predicting social tags from acoustic features on large music databases. *J. New Music Research*, 37(2):115–135.

[Bertin-Mahieux et al., 2010] Bertin-Mahieux, T., Eck, D., and Mandel, M. (2010). Automatic tagging of audio: The state-of-the-art. In Wang, W., editor, *Machine Audition: Principles, Algorithms and Systems*. IGI Publishing.

[Bertin-Mahieux et al., 2011] Bertin-Mahieux, T., Ellis, D. P., Whitman, B., and Lamere, P. (2011). The million song dataset. In *Proc. ISMIR*.

[Coviello et al., 2012] Coviello, E., Vaizman, Y., Chan, A., and Lanckriet, G. (2012). Multivariate autoregressive mixture models for music auto-tagging. In *Proc. ISMIR*.

[Flexer, 2007] Flexer, A. (2007). Musical genre classification and the artist filter effect. In *Proc. ISMIR*.

[Fu et al., 2011] Fu, Z., Lu, G., Ting, K. M., and Zhang, D. (2011). A survey of audio-based music classification and annotation. *IEEE Trans. Multimedia*, 13(2):303–319.

[Gersho and Gray, 1991] Gersho, A. and Gray, R. M. (1991). *Vector Quantization and Signal Compression*. Kluwer Academic, Norwell, MA.

[Gouyon et al., 2006] Gouyon, F., Klapuri, A., Dixon, S., Alonso, M., Tzanetakis, G., Uhle, C., and Cano, P. (2006). An experimental comparison of audio tempo induction algorithms. *IEEE Trans. Audio, Speech, Lang. Process.*, 14(5):1832–1844.

[Jensen et al., 2009] Jensen, J. H., Christensen, M. G., Ellis, D. P. W., and Jensen, S. H. (2009). Quantitative analysis of a common audio similarity measure. *IEEE Transactions on Audio, Speech and Language Processing*, 17(4):693–703.

[Lamere, 2008] Lamere, P. (2008). Social tagging and music information retrieval. *J. New Music Research*, 37(2):101–114.

[Langlois and Marques, 2009] Langlois, T. and Marques, G. (2009). A music classification method based on timbral features. In *Proc. ISMIR*.

[Law et al., 2009] Law, E., West, K., Mandel, M. I., Bay, M., and Downie, J. S. (2009). Evaluation of algorithms using games: The case of music tagging. In *Proc. ISMIR*, pages 387–392.

[Mandel and Ellis, 2008] Mandel, M. and Ellis, D. (2008). A web-based game for collecting music metadata. *J. New Music Research*, 37(2):151–165.

[Marques et al., 2011] Marques, G., Domingues, M., Langlois, T., and Gouyon, F. (2011). Three current issues in music autotagging. In *Proc. ISMIR*.

[Mauch and Ewert, 2013] Mauch, M. and Ewert, S. (2013). The audio degradation toolbox and its application to robustness evaluation. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 83–88, Curitiba, Brazil.

[Miotto et al., 2010] Miotto, R., Barrington, L., and Lanckriet, G. R. G. (2010). Improving auto-tagging by modeling semantic co-occurrences. In *Proc. ISMIR*, pages 297–302.

[Miotto and Lanckriet, 2012] Miotto, R. and Lanckriet, G. (2012). A generative context model for semantic music annotation and retrieval. *IEEE Trans. Audio, Speech, Lang. Process.*, 20(4):1096–1108.

[Nam et al., 2012] Nam, J., Herrera, J., Slaney, M., and Smith, J. (2012). Learning sparse feature representations for music annotation and retrieval. In *Proc. ISMIR*.

[Ness et al., 2009] Ness, S. R., Theocharis, A., Tzanetakis, G., and Martins, L. G. (2009). Improving automatic music tag annotation using stacked generalization of probabilistic svm outputs. In *Proc. ACM Multimedia*, pages 705–708.

[Pampalk et al., 2005] Pampalk, E., Flexer, A., and Widmer, G. (2005). Improvements of audio-based music similarity and genre classification. In *Proc. ISMIR*.

[Pan and Yang, 2010] Pan, S. and Yang, Q. (2010). A survey on transfer learning. *IEEE Trans. on Knowledge and Data Eng.*, 22(10):1345–1359.

[Panagakis et al., 2009] Panagakis, Y., Kotropoulos, C., and Arce, G. R. (2009). Music genre classification via sparse representations of auditory temporal modulations. In *Proc. European Signal Process. Conf.*

[Quionero-Candela et al., 2009] Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. (2009). *Dataset Shift in Machine Learning*. The MIT Press.

[Serra et al., 2013] Serra, X., Magas, M., Benetos, E., Chudy, M., Dixon, S., Flexer, A., Gómez, E., Gouyon, F., Herrera, P., Jordà, S., Paytuvi, O., Peeters, G., Schlüter, J., Vinet, H., and Widmer, G. (2013). *Roadmap for Music Information ReSearch*. Creative Commons BY-NC-ND 3.0 license.

[Seyerlehner et al., 2010] Seyerlehner, K., Widmer, G., Schedl, M., and Knees, P. (2010). Automatic music tag classification based on block-level features. In *Proc. Sound and Music Computing Conf.*

[Sigurdsson et al., 2006] Sigurdsson, S., Petersen, K. B., and Lehn-Schiøler, T. (2006). Mel frequency cepstral coefficients: An evaluation of robustness of mp3 encoded music. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*.

[Sordo, 2012] Sordo, M. (2012). *Semantic Annotation of Music Collections: A Computational Approach*. PhD thesis, Universitat Pompeu Fabra.

[Sturm, 2012] Sturm, B. L. (2012). Two systems for automatic music genre recognition: What are they really recognizing? In *Proc. ACM MIRUM*.

[Sturm, 2014a] Sturm, B. L. (2014a). Making explicit the formalism underlying evaluation in music information retrieval research: A look at the MIREX automatic mood classification task. In *Post-proceedings of the 2013 Computer Music Modeling and Research*.

[Sturm, 2014b] Sturm, B. L. (2014b). A simple method to determine if a music information retrieval system is a "horse". *IEEE Transactions on Multimedia*, (accepted).

[Sturm et al., 2014] Sturm, B. L., Kereliuk, C., and Pikrakis, A. (2014). A closer look at deep learning neural networks with low-level spectral periodicity features. In *Proc. Int. Workshop on Cognitive Info. Process.*

[Sturm and Noorzad, 2012] Sturm, B. L. and Noorzad, P. (2012). On automatic music genre recognition by sparse representation classification using auditory temporal modulations. In *Proc. CMMR*.

[Sugiyama et al., 2007] Sugiyama, M., Krauledat, M., and Müller, K. (2007). Covariate shift adaptation by importance weighted cross validation. *J. Mach. Learn. Res.*, 8:985–1005.

[Szegedy et al., 2014] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. (2014). Intriguing properties of neural networks. In *Proc. International Conference on Learning Representations*.

[Tingle et al., 2010] Tingle, D., Kim, Y. E., and Turnbull, D. (2010). Exploring automatic music annotation with acoustically-objective tags. In *Proc. ACM MIR*, pages 55–62.

[Turnbull et al., 2008] Turnbull, D., Barrington, L., Torres, D., and Lanckriet, G. (2008). Semantic annotation and retrieval of music and sound effects. *IEEE Trans. Audio, Speech, Lang. Process.*, 16.

[Urbano et al., 2014] Urbano, J., Bogdanov, D., Herrera, P., Gomez, E., and Serra, X. (2014). What is the effect of audio quality on the robustness of mfccs and chroma features? In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*.

[Urbano et al., 2013] Urbano, J., Schedl, M., and Serra, X. (2013). Evaluation in music information retrieval. *J. Intell. Info. Systems*, 41(3):345–369.

[van den Berg and Friedlander, 2008] van den Berg, E. and Friedlander, M. P. (2008). Probing the Pareto frontier for basis pursuit solutions. *SIAM J. on Scientific Computing*, 31(2):890–912.

[Wright et al., 2009] Wright, J., Yang, A. Y., Ganesh, A., Sastry, S. S., and Ma, Y. (2009). Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Machine Intell.*, 31(2):210–227.

[Xie et al., 2011] Xie, B., Bian, W., Tao, D., and Chordia, P. (2011). Music tagging with regularized logistic regression. In *Proc. ISMIR*, pages 711–716.

Appendices

A — Tags defining "Vocals" in CAL500 and MSD24k

```
Instrument_-_Backing_vocals
Instrument_-_Female_Lead_Vocals
Instrument_-_Male_Lead_Vocals
Vocals-Aggressive
Vocals-Altered_with_Effects
Vocals-Breathy
Vocals-Call_&_Response
Vocals-Duet
Vocals-Emotional
Vocals-Falsetto
Vocals-Gravelly
Vocals-High-pitched
Vocals-Low-pitched
```

```
Vocals-Monotone
Vocals-Rapping
Vocals-Screaming
Vocals-Spoken
Vocals-Strong
Vocals-Vocal_Harmonies
Instrument_-_Female_Lead_Vocals-Solo
Instrument_-_Male_Lead_Vocals-Solo
```

Listing 1: CAL500 tags for tag Vocals

```
a_breathy_male_lead_vocalist
a_distinctive_male_lead_vocal
a_dynamic_female_vocalist
a_dynamic_male_vocalist
a_female_vocal
a_gravelly_male_vocalist
a_laid_back_female_vocal
a_smooth_female_lead_vocal
a_smooth_male_lead_vocalist
a_vocal-centric_aesthetic
an_aggressive_male_vocalist
an_emotional_female_lead_vocal_performance
an_emotional_male_lead_vocal_performance
jazz_vocals
```

Listing 2: MSD24k tags for tag Vocals.

B — Listening test

The listening test includes 32 stimuli of 30 s each (8 stimuli with singing voice, 8 without, and their 16 transformed versions). The stimuli and one test sound example are normalized with respect to loudness. The listening test was performed online via a web-based questionnaire, written in English. The questionnaire was available online between 15th July-2nd August 2013. Few participants reported sound playback issues, consequently their responses were not included in the analyses.

Before proceeding to the experiments, participants were asked to set up the volume to a comfortable level by listening to a test sound example (not included in the stimuli). Each participant listened to the 32 stimuli and was asked to rate whether yes or no it contained a singing voice. An entire session took between 16-20 min to complete. By listening to the full list of stimuli, participants rated both conditions (original and transformed) of each stimuli. In order to control for a potential bias in ratings of the second condition heard, that would result from having previously heard the other condition, participants were assigned to one of 2 groups corresponding to a difference in presentation order: group A listened to the 16 original stimuli first and then to the 16 transformed stimuli, while group B did the opposite. Within each 16-stimuli block, the ordering of stimuli was done randomly on a subject-by-subject basis. Subjects were attributed group A or B in an alternate fashion. Participants could listen to each stimulus only once, and

they had to listen to the full duration of the stimuli before being able to listen to the next one.

A total of 254 participants took the test, of which 152 fully completed the test (79 men, 73 women, average age $\pm\ \sigma =\ 25.3y \pm 6.3$). The participants were recruited via emails, sent to international mailing lists. Participants were not paid. The following analyses are based on the 152 complete responses. There are 76 participants in both groups A and B.

Overall, the recognition of the presence of singing voice was very good, i.e. 98.1%±1.6. Considering all different conditions (original stimuli, transformed stimuli, group A, group B), and all combinations of conditions, correct recognition rates range between 97-99%.

One might raise the question whether listening to the same original and transformed stimuli successively might have implied a bias in recognition rates, i.e. artificially higher recognition rates for transformed stimuli for participants of group A, and inversely, higher recognition rates for original stimuli for participants of group B. A paired t-test was therefore conducted to compare recognition rates of singing voice presence for group A in original vs. transformed stimuli conditions. There was no significant difference in the recognition rates for original ($M = 97.5\%$, $SD = 2.5$) and transformed conditions ($M = 98.0\%$, $SD = 2.0$); $t(15) = -1.19$, $p = 0.25$. A similar test was conducted for group B. Here also, there was no significant difference in the recognition rates for transformed ($M = 98.4\%$, $SD = 1.3$) and original conditions ($M = 98.3\%$, $SD = 1.9$); $t(15) = 0.25$, $p = 0.80$. These results suggest that listening to the two conditions in a row did not imply a bias in participants recognition rates. Which therefore leads us to validate our chosen experimental design and to use the full amount of data collected in further analyses.

We performed a two-way ANOVA in order to determine whether (i) the presentation order (i.e. original version first, or alternatively transformed versions first) and, most importantly, (ii) the stimuli condition (original vs. transformed), had an effect on correct recognition of singing voice in stimuli.

The results showed no significant effect of the presentation order ($F(1, 60) = 1.59$, $p = 0.21$), hence corroborating results reported above, and no significant effect of the stimuli condition ($F(1, 60) = 0.35$, $p = 0.56$). We also found that the interaction effect between condition and presentation order was non-significant ($F(1, 60) = 0.17$, $p = 0.68$).

These results indicate that the transformation procedures do not appear to have any noticeable effect on human perception of the presence/absence of singing voice.